arXiv:1703.06452v1 [cs.CV] 19 Mar 2017

Deep Neural Networks for Semantic Segmentation of Multispectral Remote Sensing Imagery

Ronald Kemker and Christopher Kanan
Chester F. Carlson Center for Imaging Science
Rochester Institute of Technology
{rmk6217, kanan}@rit.edu

Abstract

A semantic segmentation algorithm must assign a label to every pixel in an image. Recently, semantic segmentation of RGB imagery has advanced significantly due to deep learning. Because creating datasets for semantic segmentation is laborious, these datasets tend to be significantly smaller than object recognition datasets. This makes it difficult to directly train a deep neural network for semantic segmentation, because it will be prone to overfitting. To cope with this, deep learning models typically use convolutional neural networks pre-trained on large-scale image classification datasets, which are then fine-tuned for semantic segmentation. For non-RGB imagery, this is currently not possible because large-scale labeled non-RGB datasets do not exist. In this paper, we developed two deep neural networks for semantic segmentation of multispectral remote sensing imagery. Prior to training on the target dataset, we initialize the networks with large amounts of synthetic multispectral imagery. We show that this significantly improves results on real-world remote sensing imagery, and we establish a new state-of-the-art result on the challenging Hamlin Beach State Park Dataset.

1. Introduction

Deep convolutional neural networks (DCNNs) that process high-resolution color (RGB) images have millions of parameters, e.g., VGG-16 has 138 million parameters [32]. Training these networks from scratch requires a vast amount of labeled data, otherwise they are prone to overfitting. In general, only large-scale datasets for RGB image classification, e.g., ImageNet, are suitable for fitting a DCNN's *randomly initialized* weights. To use DCNNs with smaller RGB image datasets, it has become standard practice to take a DCNN that has already been trained (i.e., pre-trained) on ImageNet and then to re-train (i.e., fine-tune) it for the new dataset or new task. With some modifications to the network, this approach has enabled state-of-the-art performance for many other computer vision tasks that use smaller datasets, including object detection [7, 27] and semantic segmentation [20, 23]. For example, to evaluate a semantic segmentation method on the RGB PASCAL VOC datasets, state-of-the-art methods use a DCNN pre-trained on ImageNet (1.28 million training images), fine-tune it for semantic segmentation on the COCO dataset (80K training images) [17], and then fine-tune it again on PACAL VOC (1,464 training images)[4, 16].

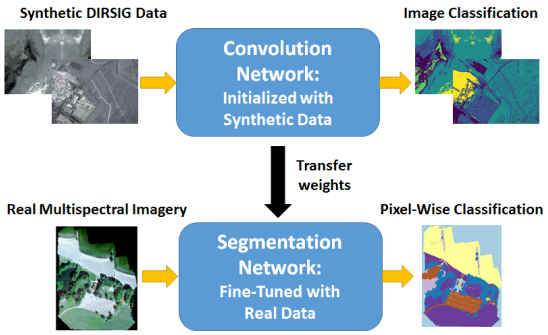

Figure 1: Our proposed model uses synthetic multispectral imagery to initialize a DCNN for semantic segmentation. This model is then fine-tuned on real imagery.

While this approach works well for RGB imagery, in the non-RGB domain, label scarcity is a far greater problem. For example, existing hyperspectral datasets for semantic segmentation of remote sensing data consist of three or fewer image mosaics. Pre-training on large hand-labeled datasets is not possible in these domains. In this paper, we explore an alternative approach: using vast quantities of automatically-labeled synthetic multispectral imagery for pre-training DCNN-based systems for semantic segmentation.

We propose to use the Digital Imaging and Remote Sensing Image Generation (DIRSIG) modeling software to gen-

erate large quantities of synthetic multispectral image data and corresponding label maps. We use DIRSIG to build a large, diverse scene model, in which we can simulate various weather and lighting conditions. We then capture synthetic aerial images of the scene with a multispectral sensor model. We use the synthetic data to pre-train a DCNN for object recognition and then fine-tune two different fully-convolutional semantic segmentation models using real imagery.

Contributions: Our paper makes three major contributions: 1) we are the first to use recent fully-convolutional neural networks for semantic segmentation with multispectral remote sensing imagery; 2) we show that pre-training these networks on synthetic DIRSIG imagery for semantic segmentation significantly increases their performance; and 3) it surpasses the previous state-of-the-art mean-class accuracy on the Hamlin Beach State Park dataset [14] by 23.6%.

2. Related Work

2.1. Semantic Segmentation with Deep Nets

Currently, state-of-the-art semantic segmentation frameworks for RGB imagery are end-to-end and consist of convolution and segmentation sub-networks. The convolution network is usually a pre-trained DCNN designed to classify images from ImageNet [20, 23, 4, 25], and current state-of-the-art performers use VGG-16 [32] or ResNet [30]. The segmentation network is appended to the convolution network and is designed to reconstruct the feature response to the same spatial dimensions as the input before assigning semantic labels. The resulting semantic segmentation network can be fine-tuned with orders of magnitude fewer training images (thousands versus millions of images) because the convolution network is already trained.

The first fully-convolutional network (FCN) designed for semantic segmentation [20] used the VGG-16 network [32], which has approximately 138 million parameters. VGG-16 was trained on ImageNet, rather than directly on much smaller semantic segmentation datasets. FCN used coarse upsampling and deconvolution in the segmentation network to classify each pixel. The net's major disadvantage was that VGG-16's 5 max-pooling layers shrunk the original image by a factor of 32, resulting in a coarse label map [20].

In [23], this FCN model was improved by building a symmetric (deconvolution) network using spatial unpooling and deconvolution layers. This increased performance when classifying small objects in the image; however, it still produced a coarse label map. As a post-processing step, the authors used a conditional random field (CRF) to sharpen the classification boundaries [15]. The major downside to this deconvolution network was that it required more memory and time to train compared to [20].

The DeepLab network [4] was built with the ResNet DCNN. To make edges sharper in the predicted label map, DeepLab uses atrous convolution in place of conventional convolution to mitigate downsampling issues. An atrous convolution filter is filled with zeros between the sample points; so although the effective size of the filter increases, the number of trainable parameters remains constant. When these filters are convolved with an image, it can preserve the original dimensions. The authors found that using atrous filters throughout the entire network was inefficient, so they used both conventional and atrous filters, reducing the image only by a factor of 8. A dense CRF was used as a post-processing step to make the predicted label map sharper.

Many of these earlier models used a CRF as a post-processing step to sharpen the label masks, but it may be better to enable the network to directly optimize itself towards creating a sharper label mask. Two recent models that did this, Sharpmask [25] and RefineNet [16], used skip-connections to incorporate image refinement into the end-to-end model. Sharpmask used a refinement module to combine features from the convolution network with the upsampled features from the segmentation network. RefineNet improved the Sharpmask model with multi-resolution fusion (MRF) to combine features at different scales, chained residual pooling (CRP) to capture background context, and residual convolutional units (RCUs) to improve end-to-end training. These components are described in more detail in Sections 3.3.1 and 3.3.2. In this paper, we adapt the Sharpmask and RefineNet models to evaluate our proposed initialization procedure.

2.2. Deep Learning for Non-RGB Sensors

Deep-learning frameworks for the semantic segmentation of multispectral and hyperspectral data have been explored by the remote sensing community; however, the lack of labeled data available for these sensor modalities have pushed researchers to embrace unsupervised feature extraction methods such as autoencoders. Deep features extracted from every pixel of the labeled data are then used with a classifier, often a support vector machine, to generate a pixel-wise classification map.

The most successful of these feature extraction methods is the stacked autoencoder [19, 18, 36, 22, 33]. An autoencoder is an unsupervised neural network that learns an efficient encoding of the training data. These autoencoders can be stacked together to learn higher feature representations.

DCNNs have been successful in RGB applications due to the large quantities of labeled data available; however, DCNNs have yet to yield state-of-the-art performance for non-RGB sensor data (e.g., [21, 18, 11]). In many cases, these DCNNS only use a single image, such as a hyperspectral image cube, to train their networks. Small quantities of training data hinder their ability to use deeper network ar-

chitectures; and in many papers, the authors fail to demonstrate that these features are transferable between datasets.

In [33], stacked sparse autoencoders (SSAE) were used to learn feature extracting filters from one hyperspectral image, and then these filters were used to classify a separate target dataset. Since the SSAE was trained on only a single dataset, the filters did not have the same discriminative power as when the SSAE was trained and used to classify the same dataset. A universal feature extractor should be trained on a large quantity of data.

In [13], the discriminative power of these unsupervised feature extracting frameworks was improved by training the model on large quantities of unlabeled hyperspectral data, which is known as self-taught learning [26]. These feature extractors yielded state-of-the-art classification performance across three benchmark datasets, showing that they can work universally well for multiple scenes.

2.3. Deep-Learning with Synthetic Data

Synthetic data has been used to increase the quantity of available training data for many applications, including object detection [24], pose estimation [5], face and handwriting recognition [35], and semantic segmentation [28]. Past work used various methods to generate large quantities of synthetic data including geometric/color transformations, 3D modeling, and virtual reality emulators.

The major upside to synthetic imagery is that its normally cheaper and easier to obtain than images that are manually annotated by humans; however, the difference in feature-space distributions, also known as the synthetic gap, makes it difficult to transfer features from synthetic to real imagery. Researchers have adopted domain adaptation techniques [3] to mitigate this phenomenon, including training autoencoders to shift the distribution of the synthetic data to the distribution of the real data [8, 35]. This can allow for a classifier to be trained using synthetic images and then make predictions on real data.

Another common form of domain adaptation, known as network fine-tuning, has been widely-adopted in the computer vision community for re-purposing DCNN frameworks for a variety of different applications, such as semantic segmentation. Portions of these DCNN frameworks were pre-trained on ImageNet and then adapted to a dataset with different class labels and feature distribution. In most cases, both these datasets contain real-imagery. In [28], the authors built a synthetic dataset using a virtual reality generator for the semantic segmentation of autonomous driving datasets. The authors combined the synthetic and real imagery to train their semantic segmentation model. In this paper, we initialized the ResNet-50 DCNN to classify synthetic imagery and then fine-tuned these weights to perform semantic segmentation on real-imagery.

Spectral Range [nm]	490-900
Spectral Bands	6
RGB Band Centers [nm]	490/550/680
NIR Band Centers [nm]	720/800/900
Spectral FWHM [nm]	10 (Bands 1-5)
	20 (Band 6)
Sensor Form Factor [pix]	1280×1024
Pixel Pitch [μm]	5.2
Focal Length [mm]	9.6
Bit Resolution	10-bit
Shutter	Global Shutter

Table 1: Specifications for the Tetracam Micro-MCA6.

3. Methods

3.1. Multispectral Sensor

The Hamlin Beach State Park dataset, used for evaluation, was collected by the Tetracam Micro-MCA6 multispectral imaging sensor. The Micro-MCA6 has six spectral bands including three color (RGB) and three near-infrared (NIR) bands. It was calibrated with an integrating sphere to properly model the sensor response in the simulated DIRSIG environment. Other technical specifications are provided in Table 1.

3.2. DIRSIG

A good DCNN initialization requires a large quantity of training data with enough spatial/spectral variability to learn discriminative feature extracting filters. There are no publicly-available ImageNet sized datasets for non-RGB sensor modalities, so we used DIRSIG to build a large synthetic labeled dataset for semantic segmentation of aerial scenes. DIRSIG is a software tool used heavily in the remote sensing industry to model imaging sensors prior to development. It can be used to simulate imaging platforms and sensor designs, including monochromatic, RGB, multispectral, hyperspectral, thermal, and light detection and ranging (LIDAR).

DIRSIG images an object using physics-based radiative transfer/propagation modeling [12, 31]. A realistic object can be defined in the DIRSIG environment with 3D geometry, textures, bi-directional reflection distribution function (BRDF), surface temperature predictions, etc. In the visible/near-infrared portion of the spectrum, DIRSIG uses MODerate resolution atmospheric TRANsmission (MODTRAN) to model atmospheric conditions in order to more accurately predict how light will propagate from the source (the sun) to the sensor. The time of day/year can be modified to change the position of the sun and moon.

The user can define any image sensor and platform configuration listed above. This includes any information about

the sensor (form-factor, spectral response, pixel-pitch, integration time), the optics (focal length), and where the sensor is looking in respect to the imaging platform. In this paper, the sensor is always looking straight down (NADIR).

A DIRSIG scene is fully customizable and can be made to be as realistic as the developer requires. The base of the scene is a 3D model and corresponding texture map of the ground (e.g. grass, sand, road, etc). Individual objects such as trees, buildings, and vehicles can be placed on top of the ground object. Each object has spatial and spectral information that will be used by DIRSIG to determine how the imaging sensor will perceive the scene. In addition to the output image, DIRSIG can output a pixel-wise classification map (semantic labels). Assuming the imaging system remains constant, the ground sample distance (GSD) can be modified by altering the height of the simulated platform.

3.3. Segmentation Networks

In this paper, we adapted the Sharpmask [25] and RefineNet [16] semantic segmentation frameworks to work for our requirements. These models were selected because they reconstruct the downsampled feature response by passing features from previous layers. These networks also learn a sharp classification mask in an end-to-end framework, resulting in higher classification performance. Our goal is to measure the benefit of pre-training with synthetic imagery on these two different networks. The convolution network for both of these models was ResNet-50 [9] with the improved network architecture scheme proposed in [10], where batch-normalization and ReLU are applied prior to each convolution. This DCNN was pre-trained using the synthetic DIRSIG imagery. Both segmentation models in this paper were developed using Theano/Keras [6].

3.3.1 Sharpmask

The Sharpmask model used for this paper is illustrated in Fig. 2. The network is broken into the convolution, bridge, and segmentation sub-networks. The convolution network is identical to the first four convolution blocks in the ResNet-50 DCNN. Since our model is being trained from scratch, we use batch normalization for regularization, whereas the original Sharpmask model did not [25].

The bridge network is a $M \times 1 \times 1$ convolution layer between the convolution and segmentation networks, where M is selected as a trade-off between performance and speed. The main goal of this network is to add some variability to the features fed into the segmentation network at refinement module #3. We use a value of $M = 512$, which worked well in preliminary experiments.

The segmentation network uses refinement modules to restore the bridge layer output to the original dimensionality of the input data. Instead of using a fully-connected

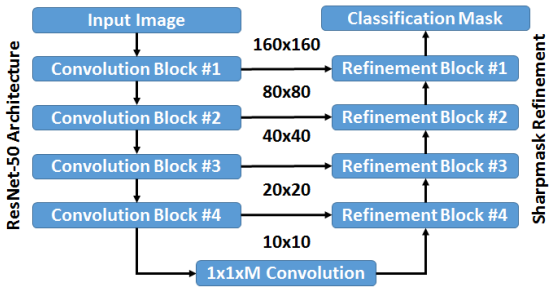

Figure 2: Our Sharpmask model. The spatial dimensions of each module's output are provided to illustrate that the refinement layer combines features from the feed-forward network (convolutional blocks) and the upsampled features in the reverse network to restore the original image dimensions. The output is a class probability matrix for each pixel.

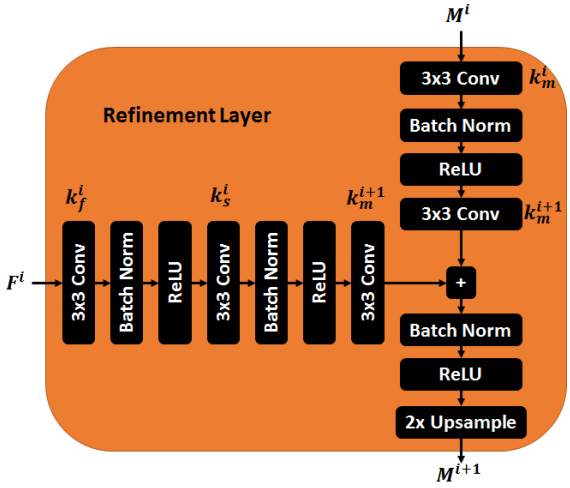

Figure 3: Sharpmask refinement with batch normalization. This layer learns mask refinement by merging features from the convolution F^i and segmentation M^i networks.

CRF, the segmentation sharpening was learned as a part of the end-to-end network. The refinement module merges low-level spatial features F_i from the convolution network with high-level semantic content in the segmentation network M_i, as illustrated in Fig. 3.

Each refinement module uses convolution and sum-wise merge layers prior to upsampling the feature response. The upsampled feature response is fed into the refinement module at the next higher dimension. The number of filters used in the i-th refinement module are given by

$$k_s^i = k_m^i = \frac{128}{2^{i-1}}, i \leq 4. \quad (1)$$

These parameters differ slightly from [25] because the

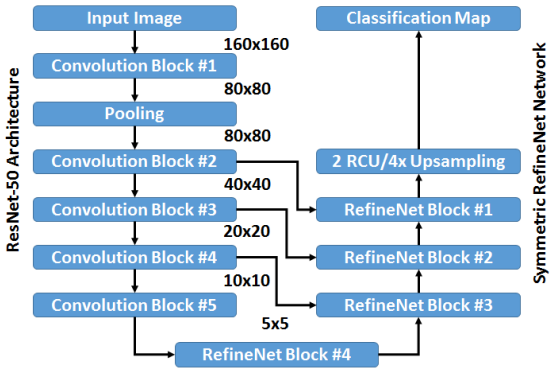

Figure 4: Network Architecture incorporating RefineNet modules for semantic segmentation.

higher dimensionality of the Tetracam images required a slightly larger model capacity.

3.3.2 RefineNet

The RefineNet model used in this paper (Fig. 4) follows the same basic structure as the Sharpmask model with a few minor changes. The refinement block in Fig. 3 is replaced with a more complex block called RefineNet (Fig. 5(a)), which is broken up into three main components: residual convolution units (RCUs), multi-resolution fusion (MRF), and chained residual pooling (CRP). In addition, our model uses batch normalization for regularization, which was not done with the original RefineNet. The convolutional net mimics the first five blocks of ResNet-50.

The RCUs (Fig. 5(b)) are used to propagate the gradient across short- and long-range connections, which makes end-to-end training more effective and efficient. The MRF is used to combine features at multiple scales, which in this paper, will be two. The CRP module (Fig. 5(c)) pools features across multiple window sizes to capture background context, which is important for discriminating classes that are spatially and spectrally similar. Fig. 5(c) uses two window sizes to illustrate how CRP works; however, our model pools features across four window sizes.

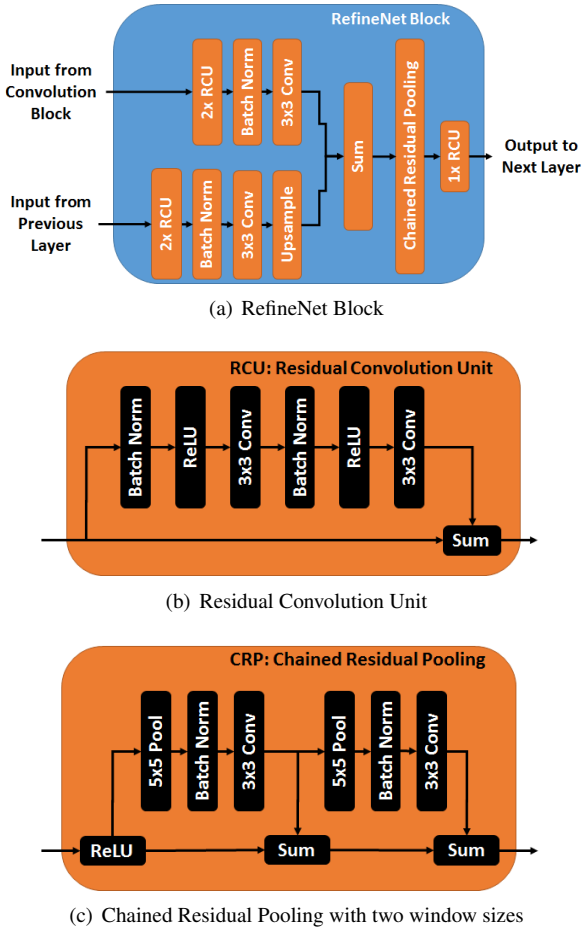

(a) RefineNet Block

(b) Residual Convolution Unit

(c) Chained Residual Pooling with two window sizes

Figure 5: RefineNet architecture.

4. Experimental Procedure

4.1. Data

4.1.1 Synthetic DIRSIG Data

We used DIRSIG to generate synthetic multispectral imagery for pre-training our semantic segmentation algorithms. We used the synthetic scene shown in Fig. 6, which resembles Trona, an unincorporated area in Southern California. It is an industrial and residential scene containing

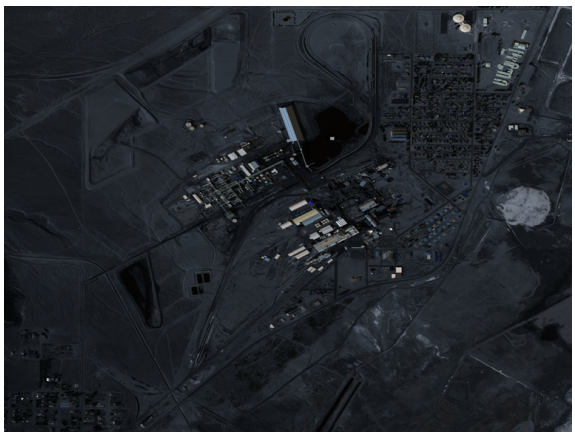

Figure 6: The DIRSIG Trona Scene.

109 labeled classes, including buildings, vehicles, swimming pools, terrain, and desert plant life.

We used calibrated radiance data captured by an integrating sphere to model the Tetracam sensor response in DIRSIG. Trona is only accurate to a GSD of 0.5 meters, and creating images with smaller GSDs produces unrealistic images due to artifacts. We generated data at multiple GSDs (0.5, 0.75, and 1 meter) to learn scale invariant features in our network. This GSD is 10 times larger than the Hamlin Beach State Park dataset.

We 'flew' the drone across the entire synthetic scene with some overlap to make sure all objects located at the edges of some images are located near the center of other images. We varied the time-of-year (summer and winter), time-of-day (morning and afternoon), and the corresponding atmospheric conditions (mid-latitude summer and winter). The final synthetic dataset contains approximately 4.7 million training and 520 thousand validation 80×80 multispectral images with their corresponding semantic labels.

4.1.2 Hamlin Beach State Park Dataset

We evaluated our deep network initialization technique on the challenging Hamlin Beach State Park dataset [14], which consists of training (9,393×5,642), validation (8,833×6,918), and test (12,446×7,654) images. The images are shown in Fig. 7. This high-resolution (4.7cm GSD) dataset was collected by the Tetracam Micro-MCA6 imaging sensor, which was flown on-board the DJI-1000 unmanned aerial system (UAS).

The Hamlin Beach State Park dataset has significant advantages over common benchmarks for multispectral and hyperspectral image segmentation in remote sensing. For many of these datasets (e.g., Indian Pines, Salinas Valley, Pavia), it is common practice to train and evaluate the algorithm on the same image, because the dataset consists of only a single image. While this might be acceptable with spectral-only features; when spatial-spectral features are used, often the training and testing folds overlap, making it difficult to assess how well an approach works in practice. While the Hamlin Beach dataset does not have many images, each one is large and very high resolution.

There are only a few benchmark datasets for multispectral imagery that correctly separate training/testing folds. The International Society for Photogrammetry and Remote Sensing (ISPRS) has two semantic segmentation challenges that uses multispectral imagery with a GSD of 5-9 cm but only six class labels [29]. The 2016 IEEE Geoscience and Remote Sensing Society (GRSS) Data Fusion Contest uses 17 class labels, but the data was resampled to 100 meter GSD [1]. The DSTL Satellite Imagery Feature Detection challenge on Kaggle has 10 classes, but the GSD of the multispectral imagery is 1.25 meters [2]. The Hamlin Beach State Park dataset has a high (4.7 cm) GSD and the 18 class labels listed in Table 2. Since a synthetic DIRSIG dataset is designed for a specific imaging system, we chose to evaluate on Hamlin Beach because it is both 1) high-resolution and 2) has a large number of classes.

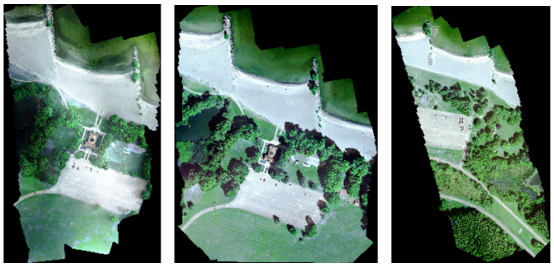

(a) Train (b) Validation (c) Test

Figure 7: RGB visualization of orthomosaic images for Hamlin Beach State Park dataset.

Previous attempts to classify the Hamlin Beach State Park dataset resulted in mediocre performance [14]. Spectral-only classification methods performed poorly because of high spatial variability found in high GSD scenes. Spatial-spectral methods performed better. The state-of-the-art performer in terms of mean-class accuracy was a spatial-spectral feature extracting framework called multi-scale independent component analysis (MICA) [13]. MICA uses unsupervised learning to build feature extracting filters from the data, and the feature responses are then fed to a SVM classifier. The problem with MICA is that it is incapable of incorporating semantic meaning of objects from outside its relatively narrow receptive field. This is where the combination of multiple convolution and downsampling operations in deep-learning frameworks, such as Sharpmask and RefineNet, can provide an advantage.

The Hamlin Beach State Park dataset is particularly challenging because its class distribution is highly unbalanced, it contains many small objects that are only present a few times in the training data, and shadows prevent some objects from being easily identified. Since all three folds are large orthomosaic images, they are broken up into smaller image patches so that the network can be trained with the memory limitations of our graphical processing units (GPUs).

4.2. Network Initialization

We compare both models with and without pre-training on synthetic DIRSIG imagery. For the pre-trained model, the ResNet-50 DCNN was initialized using 80×80 pixel patches, which is approximately 4.7 million DIRSIG training images. The network was trained using a mini-batch size of 128 and a weight-decay of 1e-4. The weights were randomly initialized from a zero-mean normal distribution.

We compute the channel-mean and standard-deviation using the entire DIRSIG training set, and these parameters were used to scale each image to zero-mean and unit-

variance. During training, the images are shuffled each epoch, and we use random horizontal and vertical flips for data augmentation. Because the class distribution for the DIRSIG data is unbalanced, we use the class-weights w_i for the entire training set to build sample (per-pixel) weights. The class-weights are scaled logarithmically using

$$w_i = \mu * \log_{10} \frac{\sum_i^N h_i}{h_i}, \quad (2)$$

where h_i is the number of pixels labeled as class i, N is the number of class labels, and μ is a tunable parameter. The network initialization stage used $\mu = 0.15$.

We optimized the network using Nadam with an initial learning rate of 2e-3. We then dropped the learning rate when the validation loss plateaued.

4.3. Network Fine-Tuning

We fine-tuned the models using the pre-trained networks to classify the Hamlin Beach State Park dataset. We randomly initialized the weights of the Sharpmask and RefineNet segmentation networks from a zero-mean normal distribution and then append them to the pre-trained ResNet DCNN. Because of the high resolution of the orthomosaics, the data is broken into 160×160 patches and fed to the segmentation frameworks. Our models are fine-tuned in two stages, using the Nadam optimizer, batch size of 32, weight decay of 1e-4, and class-weight parameter of $\mu = 0.25$.

First, the convolution network is frozen and the remaining layers are trained to adapt the weights of the segmentation network to the pre-trained weights in the convolution network. An initial learning rate of 2e-3 is used for this stage, and then it is dropped by a factor of 10 when validation loss plateaus. Second, all weights are trained using an initial learning rate of 2e-5; and again, drop it by a factor of 10 four times as the validation loss plateaued.

5. Experimental Results

Table 2 shows the classification performance of our models, with (Sim) and without (Rdm) initializing the network using synthetic data. Our networks are compared to MICA, the previous state-of-the-art method for this benchmark.

The Sharpmask-Rdm and RefineNet-Rdm networks are trained in a single-stage using the Nadam optimizer, batch-size of 32, weight decay of 1e-4, and class-weight parameter of $\mu = 0.25$. The learning rate is dropped by a factor of 10 four times as the validation loss plateaued. All of the weights are randomly initialized from a zero-mean distribution. We observed that the pre-trained models converged to a local minimum more quickly (fewer epochs) than the randomly initialized networks. Both of these models perform better on classes with more samples, but the mean-class accuracy shows that these models are not as discriminative as their counterparts that are initialized with synthetic data.

	MICA	Sharpmask		RefineNet	
		Rdm	Sim	Rdm	Sim
Road Markings	43.2	**75.9**	63.5	0.0	59.3
Tree	92.6	**94.6**	89.8	0.0	89.8
Building	1.0	0.2	**5.4**	0.0	17.0
Vehicle	47.5	38.9	20.8	0.0	**61.9**
Person	0.0	31.7	**56.9**	0.0	36.8
Lifeguard Chair	50.8	58.4	68.7	0.0	**80.6**
Picnic Table	0.0	26.5	49.3	0.0	**62.4**
Black Panel	0.0	0.0	0.0	0.0	0.0
White Panel	0.0	18.4	**54.7**	0.0	47.7
Orange Pad	66.3	**100.0**	99.6	0.0	**100.0**
Buoy	0.0	69.5	61.0	0.0	**85.8**
Rocks	66.5	**83.4**	74.4	26.5	81.2
Low Vegetation	13.3	**20.5**	1.3	0.0	1.2
Grass/Lawn	84.8	87.2	85.9	**98.8**	90.3
Sand/Beach	78.2	81.7	**94.8**	6.8	92.2
Water (Lake)	89.1	**95.3**	90.2	82.6	93.2
Water (Pond)	**3.4**	0.0	0.0	0.1	0.0
Asphalt	46.9	**97.1**	94.0	90.4	77.8
OA	74.9	**84.1**	81.3	42.0	80.4
AA	36.2	54.4	56.1	17.0	**59.8**
κ	0.683	**0.796**	0.761	0.288	0.749

Table 2: Per-class accuracies as well as overall accuracy (OA), mean-class accuracy (AA), and kappa statistic (κ) on Hamlin Beach State Park test set. The two initializations used for our Sharpmask and RefineNet models include random initialization (Rdm) and a network initialized with synthetic data (Sim). We compare our results against MICA, the previous state-of-the-art method on this dataset [14].

RefineNet-Rdm overfit to the classes with the most samples, whereas RefineNet-Sim was more discriminative with the pre-trained ResNet-50 weights. Sharpmask-Rdm did not appear to overfit. This is likely because it has fewer trainable parameters (Sharpmask and RefineNet have 11.9 and 69 million trainable parameters respectively), and the number of training samples in the Hamlin Beach State Park dataset seem to be sufficient. The pre-trained weights used in Sharpmask-Sim provided just enough assistance to boost mean-class accuracy by 1.7%. Expanding to deeper networks such as ResNet-152 would likely require the synthetic data to prevent overfitting.

Our models were unsuccessful at classifying the black panel, low-level vegetation, and the pond. The black panel has a small number of training samples; however, [14] demonstrated that simple target detection algorithms may be more appropriate for finding both wooden panels. The low-level vegetation and pond water share too many spatial-spectral characteristics with classes that have more sam-

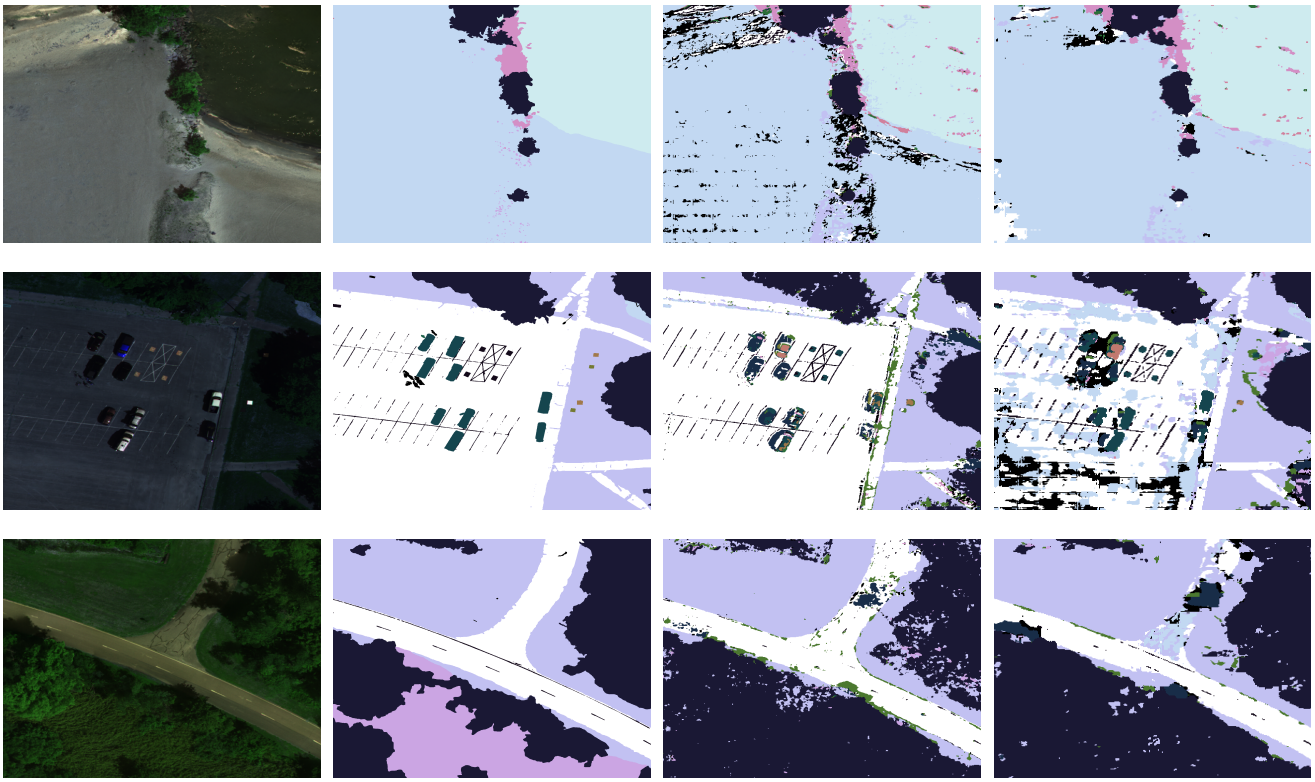

(a) Test Image (b) Ground Truth (c) Sharpmask (d) RefineNet

Figure 8: Experimental results for Sharpmask and RefineNet models.

ples. One possible solution would be to expand the input patch size, so that the network learns semantic content about neighboring pixels, e.g., the pond is surrounded by marsh and trees and the lake is surrounded by sand.

Fig. 8 shows a sample of the predictions made by our Sharpmask and RefineNet frameworks. Sharpmask does a better job at classifying the road, road markings, and vehicles; and RefineNet had fewer classification artifacts over the beach area. Both models did a good job classifying the grass and, for the most part, the lake; however, the low-level vegetation area seemed to be mis-classified as trees - which has orders of magnitude more training samples.

6. Discussion and Conclusions

In this paper, we used synthetic multispectral data to boot-strap a DCNN for semantic segmentation. The features learned from the synthetic data successfully transfered to real-world imagery. Our RefineNet-Sim model was evaluated on the Hamlin Beach State Park dataset where it beat the previous state-of-the-art mean-class accuracy by 23.6%. This work will enable remote sensing researchers to take advantage of advancements in deep-learning that were not previously available to them due to the lack of labeled data.

In the future, we hope to improve our model by 1) exploring deeper ResNet models; 2) using newer state-of-the-art convolution (ResNeXt [34]) and segmentation models; and 3) improving the inherent GSD of the DIRSIG scene; and 4) including additional diverse classes to the synthetic data. These techniques should aid the development of more discriminative frameworks that yield superior performance.

DIRSIG could be used to generate large custom datasets for other imaging modalities such as multispectral, hyperspectral, LIDAR, or a combination of all the above. These types of scenes could be developed thanks to the increase in UAS collection of remote sensing data. Our work could be adapted to these sensors. Going from multispectral to hyperspectral data will likely require modifications to our current models in order to deal with the higher dimensionality. Initializing networks using DIRSIG could improve, not only semantic segmentation, but also networks for other tasks, such as object detection in hyperspectral imagery.

Acknowledgements

We would like to thank Michael Gartley, Carl Salvaggio, and the RIT Signature Interdisciplinary Research Area, UAS Research Laboratory for their support.

References

[1] 2017 IEEE GRSS Data Fusion Contest, 2017. 6

[2] DSTL Satellite Imagery Feature Detection, 2017. 6

[3] S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. W. Vaughan. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010. 3

[4] L. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *arXiv preprint*, arXiv:1606.00915, 2016. 1, 2

[5] W. Chen, H. Wang, Y. Li, H. Su, Z. Wang, C. Tu, D. Lischinski, D. Cohen-Or, and B. Chen. Synthesizing training images for boosting human 3d pose estimation. In *3D Vision*, pages 479–488. IEEE, 2016. 3

[6] F. Chollet. Keras. https://github.com/fchollet/keras, 2015. 4

[7] R. Girshick. Fast r-cnn. In *ICCV*, 2015. 1

[8] I. Guyon, G. Dror, V. Lemaire, G. Taylor, and D. W. Aha. Unsupervised and transfer learning challenge. In *Neural Networks (IJCNN)*, pages 793–800. IEEE, 2011. 3

[9] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 4

[10] K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. In *ECCV*, 2016. 4

[11] W. Hu, Y. Huang, L. Wei, F. Zhang, and H. Li. Deep convolutional neural networks for hyperspectral image classification. *Journal of Sensors*, 2015, 2015. 2

[12] E. Ientilucci and S. Brown. Advances in wide area hyperspectral image simulation. *Proc SPIE*, 5075:110–121, 2003. 3

[13] R. Kemker and C. Kanan. Self-taught feature learning for hyperspectral image classification. *IEEE Transactions on Geoscience and Remote Sensing*, 2016. 3, 6

[14] R. Kemker, C. Salvaggio, and C. Kanan. High-resolution multispectral dataset for semantic segmentation. *arXiv preprint*, arXiv:1703.01918, 2017. 2, 6, 7

[15] P. Krähenbühl and V. Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *NIPS*, 2011. 2

[16] G. Lin, A. Milan, C. Shen, and I. D. Reid. RefineNet: Multi-path refinement networks for high-resolution semantic segmentation. In *CVPR*, 2017. 1, 2, 4

[17] T. Lin, M. Maire, S. J. Belongie, L. D. Bourdev, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014. 1

[18] Z. Lin et al. Spectral-spatial classification of hyperspectral image using autoencoders. In *Information, Communications and Signal Processing*, pages 1–5. IEEE, 2013. 2

[19] Y. Liu et al. Hyperspectral classification via deep networks and superpixel segmentation. *International Journal of Remote Sensing*, 36(13):3459–3482, 2015. 2

[20] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 1, 2

[21] F. Luus, B. Salmon, F. van den Bergh, and B. Maharaj. Multiview deep learning for land-use classification. *IEEE Geoscience and Remote Sensing Letters*, 12(12):2448–2452, 2015. 2

[22] X. Ma et al. Hyperspectral image classification via contextual deep learning. *EURASIP Journal on Image and Video Processing*, 2015(1):1–12, 2015. 2

[23] H. Noh, S. Hong, and B. Han. Learning deconvolution network for semantic segmentation. In *ICCV*, 2015. 1, 2

[24] X. Peng, B. Sun, K. Ali, and K. Saenko. Learning deep object detectors from 3d models. In *ICCV*, pages 1278–1286, 2015. 3

[25] P. H. O. Pinheiro, T. Lin, R. Collobert, and P. Dollár. Learning to refine object segments. In *ECCV*, 2016. 2, 4

[26] R. Raina, A. Battle, H. Lee, B. Packer, and A. Y. Ng. Self-taught learning: transfer learning from unlabeled data. In *ICML*. ACM, 2007. 3

[27] J. Redmon, S. K. Divvala, R. B. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016. 1

[28] G. Ros, L. Sellart, J. Materzynska, D. Vazquez, and A. M. Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *CVPR*, pages 3234–3243, 2016. 3

[29] F. Rottensteiner, G. Sohn, J. Jung, M. Gerke, C. Baillard, S. Benitez, and U. Breitkopf. The isprs benchmark on urban object classification and 3d building reconstruction. *ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci*, 1(3):293–298, 2012. 6

[30] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. S. Bernstein, A. C. Berg, and F. Li. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 115(3):211–252, 2015. 2

[31] J. Schott, S. Brown, R. Raqueo, H. Gross, and G. Robinson. An advanced synthetic image generation model and its application to multi/hyperspectral algorithm development. *Canadian Journal of Remote Sensing*, 25(2):99–111, 1999. 3

[32] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2014. 1, 2

[33] C. Tao, H. Pan, Y. Li, and Z. Zou. Unsupervised spectral–spatial feature learning with stacked sparse autoencoder for hyperspectral imagery classification. *IEEE Geoscience and Remote Sensing Letters*, 12(12):2438–2442, 2015. 2, 3

[34] S. Xie, R. B. Girshick, P. Dollár, Z. Tu, and K. He. Aggregated residual transformations for deep neural networks. In *CVPR*, 2017. 8

[35] X. Zhang, Y. Fu, S. Jiang, L. Sigal, and G. Agam. Learning from synthetic data using a stacked multichannel autoencoder. In *ICML*, pages 461–464. IEEE, 2015. 3

[36] W. Zhao, Z. Guo, J. Yue, X. Zhang, and L. Luo. On combining multiscale deep learning features for the classification of hyperspectral remote sensing imagery. *International Journal of Remote Sensing*, 36(13):3368–3379, 2015. 2